COGNITIVE PSYCHOLOGY 26, 64–101 (1994)

Below the Surface: Analogical Similarity and Retrieval Competition in Reminding

CHARLES M. WHARTON, KEITH J. HOLYOAK, PAUL E. DOWNING, TRENT E. LANGE, THOMAS D. WICKENS, AND ERIC R. MELZ

University of California at Los Angeles

While the importance of surface semantic similarities between reminding cues and memory targets has been well documented, it has been unclear whether or when human memory retrieval is influenced by structural consistency (i.e., analogical similarity), the central component of analogical mapping. The ARCS model (Analog Retrieval by Constraint Satisfaction; Thagard, Holyoak, Nelson, & Gochfeld, 1990) predicts that structural consistency, as well as competition between alternative targets in memory, will indeed influence reminding. Subjects in the experiments reported here read target passages presented in an incidental learning paradigm. After a delay, subjects read other passages and wrote down any of the previously studied texts of which they were reminded by these cues. In Experiments 1–2, related cue and target sentences contained semantically similar nouns. Cue/target consistency was manipulated by varying the case–role correspondences of these nouns. In Experiment 3, related cue and target stories contained variations of the same events, and structural consistency at the level of themes was manipulated by varying the sequencing of events. In all experiments, retrieval competition was manipulated by presenting subjects with cue sentences that were related to either a single consistent or inconsistent target (singleton condition) or both a consistent and inconsistent target (competition condition). Results indicated that both retrieval competition and structural consistency influence reminding, with the impact of the latter tending to be greater in the competition condition. The pattern of results was simulated using the ARCS model. We discuss the implications of these findings for other psychological and artificial intelligence models of memory retrieval.

Preparation of this article was supported by Contract MDA 903-89-K-0179 from the Army Research Institute, and NSF Grant DIR-9024251 to the UCLA Cognitive Science Research Program. We would like to acknowledge the various contributions of Michael Anderson, Robert Bjork, Patricia Cheng, Dedre Gentner, Douglas Medin, Mary Jo Rattermann, Bobbie Spellman, Paul Thagard, and several anonymous reviewers. We thank Rey Castillo, Vere Chappel, Christine Fax, Chris Geller, Donna Gerardi, Matt Motroni, David Thau, Sia Sipsas, and Jenny Wong for their help with data collection and analysis. Experiments 1A and 1B were reported at the Thirteenth Annual Conference of the Cognitive Science Society, Chicago, IL, August, 1991. Experiments 2 and 3 were reported at the Fourteenth Annual Conference of the Cognitive Science Society, Bloomington, IN, July, 1992. Copies of experimental materials are available from the first author upon request. Correspondence address: Charles Wharton, Department of Psychology, University of California, Los Angeles, Los Angeles, CA 90024-1563; wharton@cognet.ucla.edu.

0010-0285/94 $9.00

One of the most intriguing qualities of human memory is its capacity to allow novel experiences to bring to mind relevant prior knowledge, even though objects and events in the new situation have never been directly associated with those involved in the remembered ones. For example, a person who sees the movie *West Side Story* for the first time is likely to be reminded of the play *Romeo and Juliet,* notwithstanding the displacement of the characters in the two works over centuries and continents (Schank, 1982). This example illustrates *analogical reminding.* Both stories are analogically connected because they systematically correspond in the relationships between their actors, actions, plans, goals, and themes. *West Side Story* and *Romeo and Juliet* involve young lovers who suffer because of the disapproval of their respective social groups, causing a false report of death, in turn leading to tragedy. It is this *structural consistency* of the events and their causal relationships in the two stories that makes them analogous, rather than simply that both involve a pair of young lovers, a disapproval, a false report, and a tragedy.

It is also possible, of course, that viewing *West Side Story* could remind one of another classic musical such as *Singing In the Rain,* or of how innocent the Jets and Sharks seem compared to contemporary gangs. Not all reminding is based on analogical similarity. However, analogical reminding is of particular importance because of its hypothesized role in learning (Gick & Holyoak, 1980, 1983; Keane, 1988). New knowledge can be generated when a *source* analog is retrieved from memory and mapped to a representation of a *target* analog that is the current focus of attention. This mapping provides the basis for transfer of suitably-adapted knowledge from the source to the target representation (Gentner, 1983). For example, recognizing the analogy to *Romeo and Juliet* while watching *West Side Story* would enable one to predict that the false report of Maria's death is going to result in Tony's death.

Our aim in the present paper is to investigate the role of analogical similarity in accessing episodes in human memory. We begin by distinguishing between the different levels of abstraction at which situations can be similar, in order to more clearly define what we mean by analogical similarity. We then point out the connections between the issue of how analogical similarity influences memory access and theoretical issues in cognitive science. We follow this discussion by reviewing previous research on this topic and then deriving predictions from the ARCS model of reminding (Analogical Retrieval by Constraint Satisfaction; Thagard, Holyoak, Nelson, & Gochfeld, 1990). These predictions emphasize the need to take account of retrieval competition in investigating analogical reminding. Three experiments designed to test these predictions are reported, and the results of Experiments 2 and 3 are simulated using ARCS.

SIMILARITY AT DIFFERENT LEVELS OF ABSTRACTION

Structural consistency, the basis of analogical similarity, is inherently defined by relations among elements in a representation, rather than simply by the elements themselves. The distinction between elements and relations has been aptly characterized by Premack (1983, p. 160): ". . . [W]e may define an element as any item that is perceived as a whole, whether it be a red dot or the city of Paris." A relation holds between "two or more elements that, when presented together, are still parsed as elements (rather than being perceived as an emergent of some kind." A complex concept (e.g., Paris) may thus function as a single "chunk" or element in a representation, as in the proposition "Paris is the capital of France" (which expresses a relation between the elements "Paris" and "France"). The same concept may, in some other context, be exploded into multiple constituent elements connected by relations, as in "The Left Bank section of Paris is separated from the Right Bank by the River Seine" (which expresses a three-place relation between three constituent elements: Left Bank, Right Bank, Seine). In the same way, *West Side Story* can either be treated as a unitary conceptual element (a movie) or as a complex structured set of elements and relations (a synopsis of its plot).

When we speak of analogical similarity, we are always referring to task contexts in which the relational structure is central. For example, it is impossible to see the analogy between *West Side Story* and *Romeo and Juliet* if they are simply viewed, respectively, as a movie and a play. Rather, it is essential to represent each situation as a complex relational structure and identify structural consistencies in the correspondences between the elements of each plot.

Structural consistency provides the basis for identifying abstract similarities between situations. Gentner (1983), Halford (1992; Halford & Wilson, 1980) and Premack (1983) have outlined similar taxonomies of abstraction. In the present proposal we will distinguish three levels of abstraction, using terminology adapted from Halford (1992; Halford, Wilson, Guo, Wiles, & Stewart, in press):

Level 1: Element Similarity

Individual elements can be similar to each other by virtue of feature overlap. Psychologically, one element can be compared in isolation to another element. Relations are not involved, and hence structural consistency plays no role.

Level 2: Relational Similarity

A set of two elements may be similar to two other elements by virtue of a similar relation that holds between each pair. For example, the propositions "Paris is larger than Nice" and "The oak is bigger than the willow" can be compared to establish the correspondences: Paris ↔ oak, Nice ↔ willow. Here, structural consistency is crucial: the mapped elements need not have any element similarity, as the correspondences can be established by matching corresponding arguments of the two similar relations.

Level 3: System Similarity

System similarity is a second-order relation of similarity—similarity of relations among relations. A set of three or more elements can be similar by virtue of a similar higher-order relation that holds among the relations within each set. For example, the Aesop's "sour grapes" fable ("A fox wanted some grapes, but couldn't reach them, so he announced to his friends that the grapes were sour anyway") can be compared to the tale of a disgruntled job-seeker ("Harry hoped to get the new position of marketing manager, but was passed over, so he told his wife the job would have been boring"). The resulting correspondences are: fox ↔ Harry, grapes ↔ job, friends ↔ wife. Here structural consistency plays an even greater role than in relational similarity, in that only the higher-order relations among relations (in this case, higher-order cause–effect relations that hold between event relations) need be directly similar; once these are matched across the two stories, role correspondences can be used to map both the first-order relations and the elements.

Both relational and system similarity thus depend on structural consistency; hence we will use the term *analogical similarity* to embrace both. That is, we will use the term "analogical similarity" to refer to structural consistency as a basis for establishing correspondences between elements of two analogs.

ANALOGICAL SIMILARITY AND THEORIES OF MEMORY ACCESS

One key theoretical issue for cognitive science concerns how similarity at the higher levels of abstraction can be detected. Structural consistency depends on a solution to the ubiquitous *binding problem*—keeping track of what element plays what role, as opposed to simply what elements are involved in some way. Many major theoretical treatments of human memory retrieval assume that both cues and memory traces are represented in terms of feature vectors (e.g., Eich, 1982; Murdoch, 1982; McClelland & Rumelhart, 1985; Raajimakers & Shiffrin, 1981). Such vectors do not have any internal relational structure. They simply represent what features are present in a situation, not the roles they play. Although such models can provide simple and often compelling accounts of how element similarity operates in memory retrieval, they are unable to represent role bindings

(Barnden, in press; Ratcliff & McKoon, 1989). To represent role information, vector representations can be augmented by features that explicitly represent each possible binding of an element to a role. However, this type of solution suffers from combinatorial explosion. The problem is particularly salient in analogical comparisons that depend on system similarity. For example, to explain how the story of the disgruntled job seeker might evoke a reminding of the sour grapes fable, a vector model might postulate a preexisting feature for something like "thing that is desired but can't be obtained and hence is denigrated," and then include this feature in the *a priori* representations of both "grapes" and "job." Such a move is neither psychologically plausible nor computationally tractable as a general solution to the binding problem (see Fodor & Pylyshyn, 1988; Fodor & McLaughlin, 1990).

Notice, however, that this problem for vector models of memory retrieval would dissolve if it were established that memory access is *not* sensitive to structural consistency (i.e., analogical similarity). This might seem an implausible state of affairs. Intuitively, it certainly seems that people at least sometimes have remindings of the system-level "sour grapes" variety (e.g., Schank, 1982). And there is overwhelming empirical evidence that people are sensitive to analogical similarities when they explicitly make mappings between situations (see Gentner, 1989, for a review). However, as we will review below, psychological experiments have yet to clearly establish that structural consistency, rather than element similarity alone, plays a role in memory retrieval. The fact that people can process analogical similarity in working memory when performing mapping tasks does not imply that they can use a comparable level of similarity in the retrieval process itself. Therefore, it remains a viable possibility that human memory retrieval proceeds entirely on the basis of element similarity, without any attention to structural consistency.

An alternative possibility is that prior experimental studies of the role of structural consistency in memory retrieval have been methodologically insensitive, and hence have failed to detect a very real phenomenon. If so, then it would be necessary to postulate a more complex, binding-sensitive retrieval system, in which structured symbolic representations play a crucial role. As we will review below, computational models that have specifically addressed analogical access generally have such a character. The experiments proposed here are intended to decide between these alternative directions for theories of human memory access.

EMPIRICAL INVESTIGATIONS OF ANALOGICAL REMINDING

Given the hypothesized importance of analogical transfer in human problem solving, it is surprising that there is little firm empirical evidence that human memory retrieval is influenced by the kind of shared structural relations that in large part define the concept of analogy (Gentner, 1983). Indeed, the most robust finding in the analogy literature is that people often fail to retrieve superficially dissimilar source analogs that share substantially similar structures with a target analog (e.g., Gick & Holyoak, 1980; Keane, 1988; Ross, 1987; Seifert, McKoon, Abelson, & Ratcliff, 1986; Spencer & Weisberg, 1986). Such negative findings suggest that the process by which episodes are recalled may not be sensitive to analogical similarity, even though such similarity is often what makes retrieved information useful in problem solving.

Some problem-solving studies have shown an effect of structural consistency on retrieval when episodes in memory share similar surface elements with the current problem. For example, Ross (1989) asked subjects to solve word problems that required application of probability principles. Each probability principle had been illustrated by a single example problem. Ross varied whether the overall cover stories of the source and target problems were similar (e.g., two problems involving the IBM motor pool) or dissimilar (e.g., a motor-pool problem and a nursery-school problem). This manipulation of direct similarity of the analogs was crossed with a variation in structural consistency—the consistency with which specific elements of each analog mapped onto each other (see Gentner & Toupin, 1986, for a similar experimental design). In the consistent condition, similar types of entities mapped onto corresponding variables in relevant equations, in that people mapped to people and artifacts mapped to artifacts. In the inconsistent or "cross-mapped" condition, the role mappings involved dissimilar entities (i.e., people mapped to artifacts and artifacts to people). Ross found that inconsistent mapping impaired retrieval of the source, but only when the overall cover stories of the source and target were similar. Holyoak and Koh (1987) also found evidence that structural consistency affects retrieval when source and target analogs share some similar elements. The above studies suggest that structural consistency influences reminding, at least when element similarity is also present.

Gentner and her associates (Gentner & Landers, 1985; Rattermann & Gentner, 1987; see Gentner, Rattermann, & Forbus, 1993) conducted experiments that directly examined the role of element and system similarity in reminding. Their results were ambiguous with respect to whether system similarity influences memory retrieval, as Rattermann and Gentner (in contrast to Gentner and Landers) did not find that system similarity influences memory retrieval. Rattermann and Gentner had subjects first read a number of stories. One week later, subjects were presented with more stories and asked to write down stories from the previous session that they were reminded of by these cues. Two types of cue/target similarity were manipulated; using Halford's taxonomy, these were element similarity (i.e., similar vs dissimilar concepts referred to by concrete nouns) and system similarity (i.e., similar vs dissimilar story themes). Relational similarity (i.e., similarity of concepts referred to by verbs) was held constant at a high level between matched cue and target stories. Crossing the two types of similarity that were varied resulted in four different types of cue/target matches. Rattermann and Gentner's four conditions (using the names given to them by the investigators) can be described as follows in terms of variations in element and system similarity:

		System similarity	
		High	Low
Element similarity	High	*Literal similarity*	*Mere appearance*
	Low	*True analogy*	*False analogy*

With Rattermann and Gentner's design, a contribution of system similarity (and hence structural consistency) to analogical access would be demonstrated if reminding occurred more often for the literal-similarity condition relative to the mere-appearance condition, or for the true-analogy condition relative to the false-analogy condition. However, the only reliable differences found were advantages for the literal-similarity and mere-appearance conditions compared to the true-analogy and false-analogy conditions (recall probabilities were 0.56, 0.53, 0.12, and 0.09, respectively). The authors concluded that reminding is primarily, though not exclusively, influenced by object or "surface" similarity alone. But although element similarity was the dominant determinant of reminding in the above studies, consistent trends suggested that system similarity may also have some impact.

A few other studies have contributed evidence that high-level thematic relations can sometimes serve as a basis for reminding (Seifert et al., 1986; Johnson & Seifert, 1992), as can relatively abstract expectation failures and explanations (Gick & McGarry, 1992; Read & Cesa, 1991). Overall, however, the factors that determine whether analogical access takes place are still not well understood (e.g., McDougal, Hammond, & Seifert, 1991). The actual experimental manipulations have varied rather unsystematically from study to study, and possible structural effects (analogical similarity) have seldom been clearly differentiated from element similarity. Not all studies have yielded reliable structural effects, and positive results have produced a rather confusing picture of effects that may interact with element similarity.

The apparent weakness of analogical similarity as a contributor to reminding in some studies may be attributable to the measures that have been used to test its impact. In particular, all the above studies used designs in which at most only a *single* target[1] was strongly semantically associated to a cue (what we will term a *singleton* design). If only one item in memory has a strong semantic link to the cue, that item may tend to be retrieved regardless of whether it is structurally similar to the cue. In the present study we introduce an alternative *competition* design, in which multiple stored items are semantically related to the cue in order to test specific predictions derived from the ARCS model of reminding.

ANALOGICAL REMINDING WITHIN A CONSTRAINT-BASED RETRIEVAL SYSTEM

The ARCS model of reminding is an extension of the ACME model of analogical mapping (Holyoak & Thagard, 1989), which postulates that processing is governed by three types of constraints: direct similarity of concepts, isomorphism (i.e., structurally consistent and one-to-one mappings), and pragmatic centrality. Direct semantic similarity concerns taxonomic relations between individual concepts, such as superordination (e.g., *dog–animal*), hyponymy (e.g., *dog–cat*), and meronymy (e.g., *tail–dog*). The ARCS model includes a frame-based semantic network modeled after WordNet, an electronic lexical reference system based on a small number of core taxonomic relations (Miller, Fellbaum, Kegl, & Miller, 1988). The similarity constraint favors mappings between concepts with close taxonomic relationships. The isomorphism constraint depends crucially on case–role relations (i.e., mapped elements should fill corresponding case roles across propositions in which they appear). Pragmatic centrality concerns a preference for mappings involving elements deemed important for goal achievement.

ARCS has been tested on relatively large data bases (e.g., accessing *Romeo and Juliet* among synopses of 24 Shakespearean plays when cued with *West Side Story*) and has been used to simulate data on analogical reminding obtained in experiments by Holyoak and Koh (1987) and Rattermann and Gentner (1987) (Thagard et al., 1990). In addition, the model is consistent with other empirical evidence concerning factors that influence reminding. For example, the model predicts that an influence of structural consistency will only be observed if the cue (as elaborated by inferences) and a stored structure have at least some direct similarity

[1] There is an unfortunate mismatch between the use of the term "target" in the analogy and memory-retrieval literatures. When referring to retrieval paradigms, a cue is said to evoke one or more targets in memory. In the context of analogical reminding, the cue is typically the target analog whereas the "targets" in memory are potential source analogs. Here and elsewhere in the context of describing reminding paradigms, the term "target" refers to an item in memory that is a potential source analog.

between their constituent concepts. Without some taxonomic links, retrieval of a given stored structure will not occur because a constraint network linking it and the cue will not be formed. Ross' (1989) finding that consistency had an impact only when the overall cover stories were similar supports this prediction.

ARCS, like ACME, is a hybrid symbolic–connectionist model. Analogs are represented as symbolic structures within a semantic network; during the retrieval process, a "mapping network" is formed to enforce the constraints, and a connectionist settling process then guides retrieval of the stored structure(s) that best satisfy the constraints. More specifically, ARCS treats analogical reminding as involving a series of steps, which can be illustrated using the simplified example depicted in Fig. 1. Here two sentences, *The pastor calmed the businessman* (*T1*) and *The executive soothed the priest* (*T2*) are stored in long-term memory, and the sentence *The rabbi reassured the chairman* (*C1*) serves as a retrieval cue.

First, linguistic comprehension processes (which ARCS does not model) will produce an elaborated semantic representation of the cue. Comprehension may include the generation of additional propositions via application of inference rules and resolution of lexical and/or syntactic ambiguities. Second, the concepts in the cue will serve as probes into long-term memory. Taxonomic relations (not illustrated in Fig. 1) will provide retrieval pathways that connect similar concepts. For the example in Fig. 1, each concept in the cue with make contact with a hyponym in each of the stored sentences (e.g., *chairman* will be connected to *businessman* and *executive, rabbi* to *pastor* and *priest,* and *reassure* to *calm* and *soothe*).

Using information about the direct similarity of predicates, a constraint network is created that includes "mapping units" that represent possible correspondences between similar predicates (e.g., *reassure* = *calm*), propositions based on these predicates (e.g., *C1* = *T1*), objects that serve as arguments of the propositions (e.g., *obj-rabbi* = *obj-pastor*), and larger episodic structures (e.g., stories). (The simple example in Fig. 1 involves single-sentence structures.) Note that only stored structures linked to the cue by similar concepts will be potentially retrieved.

Excitatory and inhibitory connections between units are then created to enforce the constraints postulated by the model. To implement the similarity constraint, mapping units for predicates that are similar in meaning (e.g., hyponyms, as in *rabbi* = *pastor*) receive excitation from a special semantic unit (always clamped on). To implement the structural consistency constraint, excitatory connections are constructed between mapping units based on a particular proposition mapping. Thus the possible mapping *C1* = *T1* has excitatory connections to the corresponding predicate mapping, *reassure* = *calm,* and to the corresponding mappings

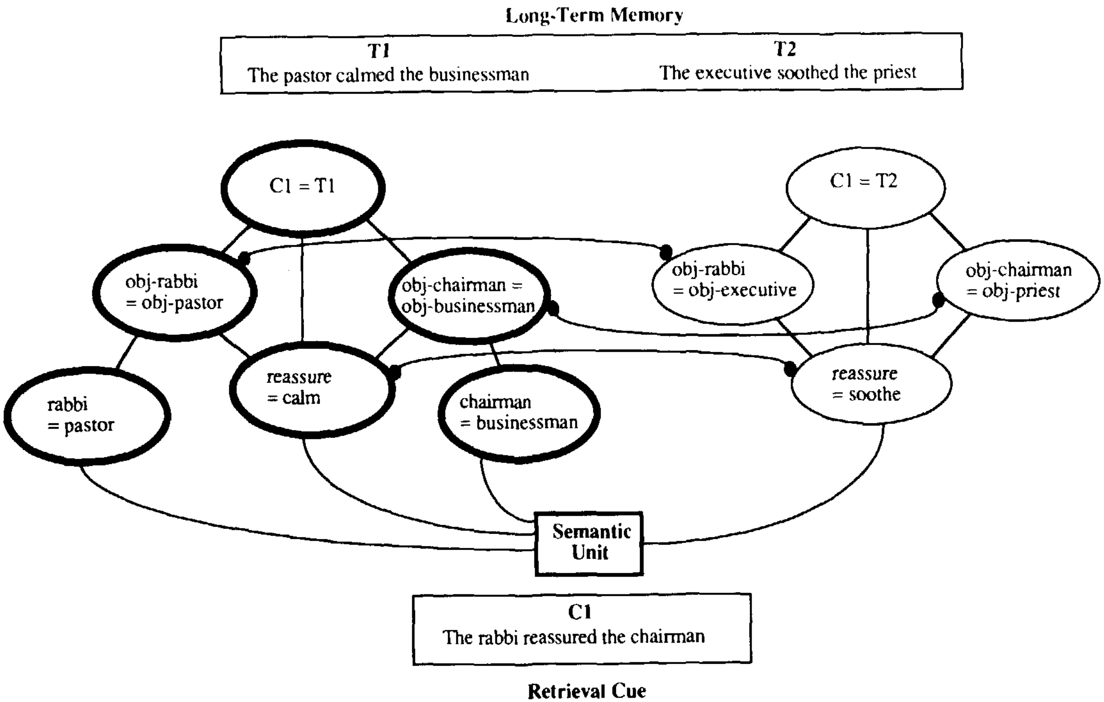

FIG. 1. Illustration of an ARCS network in which a structurally consistent and a structurally inconsistent target are competing to be retrieved. L with solid circles are inhibitory connections between mapping units.

of arguments, *obj-rabbi* = *obj-pastor* and *obj-chairman* = *obj-businessman*. To encourage mappings to be one-to-one (the other aspect of the isomorphism constraint), inhibitory connections are placed between units for alternative mappings (e.g., between *obj-rabbi* = *obj-pastor* and *obj-rabbi* = *obj-executive*). Finally, to implement the constraint of pragmatic centrality, another special unit provides further excitation to any mapping units involving elements deemed especially important (by virtue, for example, of goal relevance). (Figure 1 does not illustrate the pragmatic-centrality constraint, which is outside the scope of the present paper.)

Once the constraint-satisfaction network has been formed, a standard connectionist settling procedure is used to find a global activation state that yields a local maximum of a "goodness" measure. Those mapping units that reach a relatively high activation level (indicated by darker borders in Fig. 1) are candidates for retrieval. In particular, the most active mapping unit involving the cue, *C1*, determines the stored structure most likely to be retrieved. For the example in Fig. 1, the unit *C1* = *T1* will be favored over *C1* = *T2*. Even though the two stored sentences have equal (and high) taxonomic similarity of their individual concepts to those in the cue, the structural relationships favor the mapping *C1* = *T1*. The predicate-mapping units *rabbi* = *pastor* and *chairman* = *businessman* are supported by semantic similarity; these units will in turn support the corresponding object–mapping units, *obj-rabbi* = *obj-pastor* and *obj-chairman* = *obj-businessman*, which will then dominate their rivals, such as *obj-rabbi* = *obj-executive* and *obj-chairman* = *obj-priest*. As a consequence, the unit *C1* = *T1* will receive greater support from its corresponding argument mappings than will *C1* = *T2*, so that the former will dominate the latter. ARCS would thus predict that *C1* is more likely to trigger reminding of *T1* than of *T2*.

As the above example illustrates, retrieval in ARCS is fundamentally competitive, in that evidence favoring access to one stored structure serves to reduce (via inhibitory links) the likelihood of retrieving other stored structures. It has long been known that learning multiple associations to a common cue leads to retrieval interference (e.g., Anderson, 1974; Barnes & Underwood, 1959), and it has been argued that some forms of retrieval inhibition are actually adaptive (Anderson & Milson, 1989; Bjork, 1989). ARCS' constraint–satisfaction algorithm naturally generates competitive effects even in reminding situations in which the cue was never directly associated with any stored structures, in which case access depends on preexisting taxonomic relations in semantic memory.

Several qualitative predictions concerning determinants of analogical reminding can be derived from ARCS. First, ARCS predicts that targets that are structurally consistent with the cue will have a retrieval advantage over targets that are not. Second, because retrieval in ARCS is fundamentally competitive, it follows that if a cue has semantic links to multiple structures in memory, then reminding of these targets will be impaired relative to when each is the sole semantically related target in memory. Third, ARCS predicts that structural consistency will interact with competition. There will be high probability of retrieving a structurally inconsistent target if it is the only semantically related structure in memory; but there will be a diminished probability of retrieving such a target if it is competing with another semantically related target that is structurally consistent with the cue. Thus for the example in Fig. 1, the structurally consistent *T1* will dominate over *T2* when both are in memory, so that retrieval of *T1* will be relatively good despite the presence of *T2*. In contrast, *T2* might be retrieved quite well if it were the only semantically related structure in memory, in which case it will be the sole candidate for retrieval. However, as the loser in competition with *T1*, retrieval of *T2* will suffer more severely in the competitive situation. ARCS thus predicts that a competition design will prove much more sensitive to the impact of structural consistency on reminding than is a singleton design of the sort used in previous studies. Finally (also reflecting the influence of retrieval competition), ARCS predicts that relatively unrelated items in memory are more likely to be retrieved when there is no target in memory that is highly similar to the cue.

All of the present experiments were designed to test these predictions. In Experiments 1–2, we varied the relational similarity between cues and targets while holding element similarity constant at a relatively high level (concepts that were cohyponyms of a common superordinate, such as *pastor* and *rabbi*). In Experiment 3, we varied the system similarity between cues and targets while holding relational similarity constant at a moderate level. Analogous and disanalogous targets were based on different story themes, but were both composed from relatively similar events (e.g., trying out for a band, interviewing for a job).

EXPERIMENT 1A

Subjects read a series of simple (2- to 3-sentence) texts in the context of various incidental tasks designed to ensure semantic processing. After a brief delay, subjects were given a series of single subject–verb–object (SVO) cues and asked to write any of the previous texts of which they were reminded by each cue sentence. An example of lexically related target and cue passages is shown in Table 1. One text contains a consistent target sentence, *The pastor calmed the businessman,* in which nouns fill case roles parallel to their respective hyponyms in the cue sentence, *The rabbi reassured the chairman.* The other text contains an

TABLE 1
Structurally Consistent and Inconsistent Target Passages with Related Cue Sentence

Consistent target	
	Having just been fired from a high level job, he decided to go to his church for counseling. *The pastor calmed the businessman.*
Inconsistent target	
	The church was having trouble approaching local corporations for contributions to the shelter. *The executive soothed the priest.*
Cue sentence	
	The rabbi reassured the chairman.

inconsistent target sentence, *The executive soothed the priest,* in which noun hyponym case roles are reversed between the target sentence and cue.[2]

Retrieval cues were lexically associated to both a consistent and an inconsistent target sentence (competition condition), a consistent target sentence (consistent singleton condition), an inconsistent target sentence (inconsistent singleton condition), or no target sentence (unrelated condition). A diagram of this design is shown in Fig. 2. Besides target sentences, target passages also contained sentences (*surround sentences*) that served to elaborate the scene being described by the target sentences (e.g., "Having just been fired from a high level job . . ."). We analyzed surround sentence recall separately from target sentence recall in order to see if ARCS predictions were supported even for components of target passages that had no direct association with cue passages. Based on ARCS, we predicted that (1) consistent passages would be recalled more often than inconsistent passages, (2) both target types would be recalled less often in the competition than the singleton condition, and (3) the difference between reminding for consistent and inconsistent targets would be greater in the competition than in the singleton condition.

Method

Materials. Materials consisted from 24 sets of texts, each set consisting of two target texts and one cue sentence. Each text consisted of a target sentence and one or two accompanying "surround" sentences intended to make the texts more meaningful. Each target

[2] Although for brevity we will refer to the two conditions as "consistent" versus "inconsistent," it is important to note that in terms of the ARCS model the variation is actually one of agreement versus disagreement between the constraint of structural consistency and that of semantic similarity. That is, in the consistent condition both constraints converge on the same mapping, whereas in the inconsistent condition the pressure to maintain structural consistency of the mapping is pitted against the pressure toward mapping similar objects.

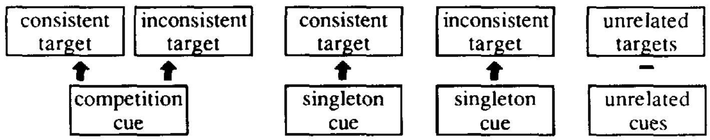

FIG. 2. Basic design of present experiments.

sentence was matched to the set's cue sentence. Matched target and cue sentences shared two sets of associated nouns (e.g., *pastor, priest, rabbi; businessman, executive, chairman*) and a set of associated verbs (e.g., *calm, soothe, reassure*). Nouns within sets were chosen so that they would jointly make sense in either the object or subject position. Within each text pair, verbs were randomly assigned to pairs of target nouns, after which each target sentence was randomly assigned to one of the two surround texts. In order to avoid confounding cue/target consistency with surface order, an equal number of active and passive cue sentences and target sentences were constructed in each condition. For each set of materials, random assignment was used to decide whether each cue sentence and each target sentence would be active or passive.

Materials were presented in booklets, one item per page. The pages of each booklet were randomly ordered. Separate booklets were created for the targets and cues. The target booklets contained 12 target texts: two competitor pairs, two consistent singletons, two inconsistent singletons, and four unrelated singletons (i.e., lexically unrelated to any presented cue). The cue booklets included 10 cue texts: two competitor cues, two consistent singletons, two inconsistent singletons, and four unrelated singletons (see diagram of experimental design in Fig. 2). Each page displayed three 6-point Likert scales for ratings of the item's plausibility, meaningfulness, and imagability. To ensure that sentences from every theme group were used in every condition, a Latin-square design was used to create 12 different configurations of cues and targets. Each configuration was administered to three subjects. Separate items were used for the four unrelated targets and four unrelated cues.

In order to determine whether people are sensitive to our consistency manipulation when retrieval is *not* required, an initial group of 96 undergraduates attending the University of California at Los Angeles (UCLA) completed brief questionnaires designed to assess the perceived similarity of the target texts and the cue sentences. Paired target and cue passages appeared in consistent, inconsistent, and unrelated conditions. Subjects were asked to rate ". . . how similar are the scenes being described" on a 6-point Likert scale (range: 1, not similar at all, to 6, identical). Consistent cue/target pairs were rated as more similar than inconsistent cue/target pairs (3.98 vs 2.75), min $F'(1,100) = 28.04$, $p < .001$; inconsistent cue/target pairs were rated as more similar than unrelated cue/target pairs (1.59), min $F'(1,75) = 36.10$, $p < .001$. (All analyses of variance reported here used the min F' statistic, Clark, 1973, to allow simultaneous generalization over both subject and material variability.)

Subjects and procedure. Subjects were 36 UCLA undergraduates who were paid for their participation. Participants were tested either individually or in small groups and were read instructions aloud by the experimenter. Subjects were told that they would be asked to read and rate a series of stories on a 6-point scale for plausibility, meaningfulness, and imagability. For all scales, a 1 represented the lowest rating (e.g., implausible) and a 6 represented the highest rating (e.g., highly imagable). Rating attributes were defined to subjects as (a) plausibility: how realistic a story was, given that it was supposed to take place "in the real world"; (b) meaningfulness: how easy a story was to understand; and (c) imagability: how easy a story was to visualize mentally. Subjects were informed that the experimenter was writing a computer simulation of how people represented story scenes as a function of these attributes, and so they should make their ratings as accurate as possible. Subjects were asked to go through the booklet three times, each time making only one attribute rating per text. Before each pass through the booklet, the experimenter provided subjects with a specific definition of each rated attribute.

For the first pass through the booklet, subjects were given 30 s to read and rate each text. For the second and third passes through the booklet, subjects were given 20 s to read and rate each text. Subjects were told that they would be signaled when to turn each page. After completion of the third pass, subjects participated in a 5-min logical reasoning experiment as a distractor task. Subjects were then informed that the computer simulation was also a model of how people were reminded of stories. Participants were told that they would now read another booklet of sentences, some of which were similar to one or more of the stories that they had read in the first part of the experiment. They were instructed to make one pass through the new booklet. Subjects were told that as they went through the booklet, they were to rate each sentence for plausibility, meaningfulness, and imagability and then write down as accurately as they could what they were reminded of from the previous passages they had rated. It was emphasized that some sentences might not remind them of any story, while others might remind them of several stories, in which case they were to recall as much as they could of each. Subjects were told that while the sentence ratings were of interest, the story or stories they were reminded of while they made these ratings were of greatest importance. No time limit was imposed, but subjects were told that they should not turn back to a page once they turned that page even if they were later reminded of something. The experimental session lasted approximately 30 min.

Results and Discussion

Overview of all data analyses. Access to target and surround sentences (i.e., target passage sentences not lexically associated to cue sentences) were scored separately. For each separate attempt at recalling a text (i.e., what a subject considered a single story), access credit was given to whichever studied text had content words recalled. If content words from two texts were included, credit was given to the text that contributed the most content words. (A single recall attempt never included content words from more than two texts.) When a retrieval attempt produced equal numbers of content words from two texts, or when both texts were recalled as separate texts, access credit was given to each. Thus if subjects wrote down content words from two texts in *separate* retrieval attempts in response to a single cue, access credit was given for each. Because synonym substitutions were easily confused with interchanges of hyponyms across paired texts, literal recall was required in scoring content words.

A desirable feature of this access measure is that it is not itself sensitive to the structure of the accessed text. A measure that required accurate recall of case roles might produce an artifactual advantage for the structurally consistent target cues. Suppose a subject who was cued with *The rabbi reassured the chairman* was uncertain of the case–role assignments in the text. A simple strategy would be to match the assignments in the recalled text to those in the cue. This strategy would produce apparent correct recall of case roles if the cue was consistent with the target (*The* *pastor calmed the businessman*) but an incorrect role reversal if the cue were inconsistent with the target (*The priest soothed the executive*).

As mentioned above, all tests of mean differences between conditions were analyzed using the min F' statistic, which required computing separate analyses of variance (ANOVA) for subjects and items. For the subject ANOVAs, materials configuration (12 levels, Experiment 1A; 24 levels, Experiments 1B and 2; 14 levels, Experiment 3) was included as a fixed between-subjects factor. For the item ANOVAs, it was important to remove variance due to subject differences from the error terms. There were only three subjects per item per condition in Experiment 1A (and also in later experiments). By necessity, the assignment of particular items to conditions had to vary across subjects. Accordingly, each subject's grand mean recall score was subtracted from the subject's raw score in each condition before item recall means in each condition were calculated. Also, in Experiments 1A and 1B, target passages from each material group were presented in either consistent competitor and inconsistent singleton conditions or in inconsistent competitor and consistent singleton conditions, so that targets from each theme set were used in these four conditions. Thus, we treated each of the 24 materials groups as a unit in the item ANOVAs (i.e., 24 observations). In Experiment 2, all 48 possible target sentences appeared in all conditions, and so there were 48 observations in the item ANOVAs. In Experiment 3, examples of all 28 story themes appeared in all conditions, and so there were 28 observations in the item ANOVAs.

Reminding analyses. Figure 3 presents the proportion of related texts accessed in response to cues, plotted separately for target and surround sentences. ARCS' predictions were supported: Consistent target sentences were recalled more frequently than inconsistent targets, min

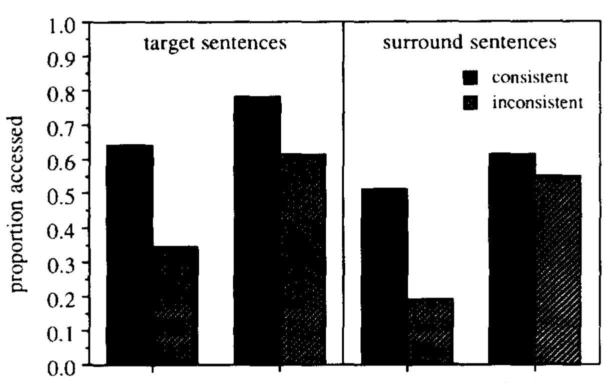

FIG. 3. Reminding for target and surround sentences in Experiment 1A.

$F'(1,42) = 9.27$, $p < .01$. Although a larger effect of consistency was obtained in the competitor condition than in the singleton condition, the interaction between consistency and competition was not significant, min $F' < 1$. Singleton target sentences were recalled more often than competitor target sentences, min $F'(1,34) = 10.95$, $p < .01$. The frequency of retrieving at least one related target (consistent or inconsistent) in the competitor condition was 85%, compared to 78% in the consistent singleton condition and 61% in the inconsistent singleton condition. In the competitor condition, frequency of accessing *both* a consistent target and an inconsistent target was 14%.

Consistent surround sentences were recalled more frequently than inconsistent surround sentences, min $F'(1,38) = 6.56$, $p < .05$. While the difference between consistent and inconsistent surrounds in the competitor condition was considerably greater than in the singleton condition, this interaction was not significant, min $F'(1,39) = 2.54$, $p > .10$. However, consistent surround sentences were recalled more often than inconsistent surround sentences in the competition condition, min $F'(1,39) = 8.57$, $p < .01$, but not in the singleton condition, min $F' < 1$. Singleton surround sentences were recalled more frequently than competitor surround sentences, min $F'(1,42) = 9.27$, $p < .01$. Frequency of accessing at least one consistent surround or inconsistent surround was 64% in the competitor condition, compared to 61% in the consistent singleton and 54% in the inconsistent singleton conditions; frequency of retrieving both a consistent surround and an inconsistent surround was 8%.

Other aspects of the reminding data support the hypothesis that reminding is characterized by retrieval competition among all targets in memory. This systematic retrieval competition is shown by the pattern of reminding of unrelated items across conditions. Recall of unrelated target sentences increased monotonically as the strength of association of other items in memory to the cue decreased (unrelated target sentence recall: competition condition, 1%; consistent singleton condition, 3%; inconsistent singleton condition, 7%; unrelated-cue condition, 17%). There was thus a low rate of reminding for unrelated targets with related-target cues because of strong competition from the texts specifically related to the target cues. In contrast, there was an increased rate of reminding for unrelated targets with unrelated cues, because there was no retrieval competition from target sentences related to these cues.[3] A similar pattern was shown with surround sentences.

[3] This pattern of reminding for unrelated targets can be analysed in a simple 3-parameter statistical model that also applies to the remindings for consistent and inconsistent targets (Wharton et al., 1991). This analysis supports the claim that reminding was increased by structural consistency and decreased by retrieval competition in an orderly fashion across all conditions. Similar results were obtained in fitting all other experiments reported here.

To explore access to the case-role structure of recalled texts, we also scored whether or not each retrieved target noun was recalled in its correct case role (agent or object). Target nouns were recalled in their correct case roles in 95% of trials in which they were retrieved. Because so few nouns were recalled in incorrect case roles, more detailed analyses were not warranted. These results indicate that although our primary access measure does not depend on retrieval of correct case roles, such information was almost always available whenever a target sentence was accessed.

EXPERIMENT 1B

Experiment 1B was performed to provide a replication of the results of Experiment 1A.

Method

Experiment 1B included the four reminding conditions tested in Experiment 1A.[4] In order to rule out the possibility that the effects found in Experiment 1A were due to chance randomization of our text materials, we used the same basic design as before but reconstructed our set of materials for Experiment 1B. Within each text pair, verbs were again randomly assigned to pairs of target nouns, after which each target sentence was randomly assigned to one of the two surround texts. For each set of materials, random assignment was again used to decide whether each cue sentence and each target sentence would be active or passive. A new order was randomly determined for the sequence in which materials rotated through experimental conditions. To ensure that every sentence was used in every condition, a Latin-square design was used to create 24 separate configurations of targets and cues, each administered to three subjects. The 24 materials configurations were divided into two equal subsets, each seen by 36 subjects.

Subjects were 72 UCLA undergraduates who participated in order to meet a course requirement for an introductory psychology course. In all other respects, the methodology of Experiment 1B was identical to that of Experiment 1A.

Results and Discussion

Figure 4 presents the proportion of related texts accessed in response to target cues, plotted separately for target and surround sentences. The findings of Experiment 1A were basically replicated: Consistent target sentences were recalled more frequently than inconsistent target sentences. This difference was only marginally significant with min $F'(1,54) = 3.56$, $p < .10$, but both subject and item ANOVAs were significant, $F(1,48) = 8.53$, $p < .01$, $F(1,23) = 6.10$, $p < .05$, respectively. As in Experiment 1A, while there was a trend toward a larger effect of consistency in the competitor condition than in the singleton condition, this interaction was not significant, min $F'(1,39) = 1.19$, $p > .10$. Singleton

[4] Three additional conditions involving cuing of surround sentences were also run in Experiment 1B, but these will not be described here.

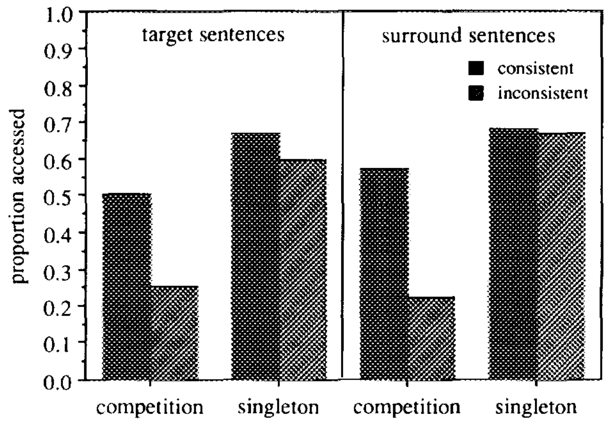

FIG. 4. Reminding for target and surround sentences in Experiment 1B.

target sentences were recalled more often than competitor target sentences, min $F'(1,68) = 20.15, p < .001$. The frequency of reminding for at least one consistent or inconsistent target sentence in the competitor condition was 68% compared to 67% in the consistent singleton condition and 60% in the inconsistent singleton condition; frequency of recalling both a consistent and an inconsistent target sentence was 7%.

As in Experiment 1A, consistent surround sentences were recalled more frequently than inconsistent surround sentences, min $F'(1,48) = 4.33, p < .05$. The interaction between consistency of mapping and target competition was significant, min $F'(1,47) = 4.87, p < .05$. Consistent surround sentences were recalled more often than inconsistent surround sentences in the competition condition, min $F'(1,43) = 6.60, p < .05$, but not in the singleton condition, min $F' < 1$. Singleton surround sentences were recalled more often than competitor surround sentences, min $F'(1,53) = 19.99, p < .001$. Overall frequency of accessing surround sentences in the competitor condition was 74% compared to 68% in the consistent singleton and 67% in the inconsistent singleton conditions. Frequency of retrieving both consistent target and inconsistent surround sentences in the competitor condition was 6%.

As in Experiment 1A, the percentage of unrelated target sentences recalled increased monotonically from the competition condition to the consistent singleton condition to the inconsistent singleton condition to the unrelated-cue condition (0, 1, 7, and 21%). Only in the absence of competition from semantically related targets (i.e., in the unrelated-cue condition) was there a substantial probability of recalling an unrelated target.

Subjects sometimes mixed content words from two texts within a single recall attempt, in which case our basic access measure gave credit to only the dominant text. Over both Experiments 1A and 1B, blended recall was more likely to occur for target than for surround sentences (26 vs 7% of trials), min $F'(1,115) = 7.34$, $p < .01$, reflecting the fact that the paired target sentences were semantically related, whereas the paired surrounds were not. For comparison, we also tabulated the frequency with which words from paired texts were inserted in the recall of singleton texts, when the paired texts had not been studied. Such errors provide a baseline estimate of the probability of substituting related hyponyms during retrieval (e.g., substituting *priest* for *pastor*). Such intrusions occurred for 7% of singleton targets and 0% of singleton surrounds. In the competitor condition, exposure to one text in a pair may have primed words that subsequently intruded during attempts to recall hyponyms that occurred in the other text in the pair.

EXPERIMENT 2

In Experiments 1A and 1B, target material scene descriptions consisted of two or three sentences. It is likely that these elaborated contexts encouraged subjects to make inferences that augmented their text representations and so increased the probability that consistent targets would be recalled. The passages in Table 1 can be used to illustrate how sentential context might make readers add additional knowledge to their text representations. For the consistent target in Table 1, it is natural to infer from the surround statement, *Having just been fired from a high-level job . . .* , that the businessman went to the pastor about a personal problem, and that the priest helped the businessman to deal with this problem. For the inconsistent target, the surround statement, *The church was having trouble approaching local corporations for contributions . . .* might lead to the inference that the priest went to the executive about a financial problem, and that the executive helped the priest to deal with it. The cue sentence, *The rabbi reassured the chairman,* although not itself embedded in a text context, seems more likely to elicit inferences that parallel those generated when reading the consistent target, *The pastor calmed the businessman,* than those generated when reading the inconsistent target, *The executive soothed the priest.* Thus, additional inferences triggered by the surrounding text context should increase semantic and structural overlap between subjects' representations of the cue sentence and of the consistent target, relative to that of the cue with the inconsistent target.

It is possible that the inferences that readers create from multiple-sentence scene descriptions are necessary to obtain effects of structural consistency in reminding. This view contrasts with ARCS' prediction that targets with consistent object mappings will be recalled more frequently than targets with inconsistent object mappings even in the absence of differential inferences. In order to test these contrasting predictions, single sentence targets *without* any surrounding text context (e.g., *The pastor calmed the businessman*) were presented to subjects during initial encoding.

Method

Except where noted, the method was same as in Experiments 1A and 1B. Materials from Experiment 1B were used except that target passages were not studied with their surround sentences. All materials were rotated through all conditions including the unrelated condition.

We collected subjects' similarity ratings in order to explore the effect of removing surround sentences from our target passages. Subjects were 96 UCLA undergraduates. In spite of the absence of the elaborating context provided by surround sentences, consistent cues and targets were rated as being more similar than inconsistent cues and targets (4.10 vs 3.12), min $F'(1,56) = 19.76$, $p < .001$. Inconsistent cues and targets were rated as being more similar than unrelated cues and targets (2.07), min $F'(1,51) = 19.11$, $p < .001$.

Subjects in the reminding task were 72 UCLA undergraduates who participated either for pay or in order to meet a requirement for an introductory psychology course. Because each target in Experiment 3 was only one sentence long, we allowed subjects 20 s (instead of 30 s) for their first rating pass through the booklets.

Results and Discussion

Figure 5 shows the proportion of target sentences accessed in response to target cues for Experiment 2. As in the previous experiments, structurally consistent targets were recalled more frequently than structurally inconsistent targets, min $F'(1,93) = 5.76$, $p < .05$. The interaction between retrieval competition and structural consistency was not significant in either the subject or the item analyses, both $F < 1$. Singleton targets were recalled more often than competitor targets, min $F'(1,92) = 33.39$, $p < .001$. Frequency of retrieving at least one related target in the compet-

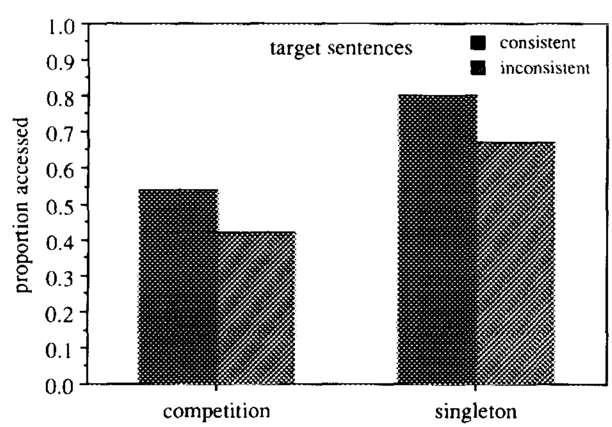

FIG. 5. Reminding for target sentences in Experiment 2.

itor condition was 86% compared to 80% in the consistent singleton condition and 67% in the inconsistent singleton condition. Access to both a consistent and an inconsistent target sentence occurred in 10% of retrievals in the competitor condition. As in the previous experiments, remindings of unrelated targets increased monotonically from the competition to the singleton consistent to the singleton inconsistent to the unrelated-cue condition (2, 6, 7, and 23%). These differences again suggest the ubiquitous influence of competitive influences on probability of reminding.

The results of Experiment 2 clearly indicate that the surround sentences used in Experiments 1A and 1B are not necessary to show the effect of structural consistency on reminding. As in the previous two experiments, structurally consistent targets were recalled more often than inconsistent targets. Except for the additional surround sentences, Experiments 1A and 1B used the same target items as in Experiment 2. While cross-experiment comparisons have to be treated with caution, Experiments 1A and 1B can be compared to Experiment 2 to explore the effect of the surround sentences. The difference between overall retrieval probabilities for consistent versus inconsistent target sentences was somewhat smaller in Experiment 2 than in Experiment 1A or Experiment 1B (12 vs. 23 and 17%, respectively). This pattern lends some support to the view that surround sentences encourage subjects to make more inferences about the target sentences, and hence increase the effective difference in overall similarity between the consistent and inconsistent targets.

Unlike the previous two experiments, the effect of structural consistency was not larger in the competition condition than in the singleton condition. One reason may be that the overall thematic differences between consistent and inconsistent target sentences are reduced when there is no surrounding discourse context. Accordingly, the competition between analogous and disanalogous targets may have been more nearly equal than in the previous two experiments.

EXPERIMENT 3

Experiments 1–2 manipulated relational similarity while holding element similarity constant at a high level. In order to determine whether structural consistency can also influence reminding when the similarity relationship is more abstract, for Experiment 3 we constructed a set of stories that allowed similarity at the system level to be varied. An example is shown in Table 2. Here, the same underlying set of types of events (e.g., employment interview, shopping) is used to construct two different story themes, which might be roughly glossed as *counting your chickens before they're hatched* (theme 1), and *finding desperately-needed employment* (theme 2). Themes depend on higher-order relations between relations and hence involve system similarity. No content words were used in

TABLE 2
Example of a Materials Set Used in Experiment 3 (after Seifert et al., 1986)

Set A, Theme 1: Ernie was really encouraged about his interview for a security guard at the new factory in town. He thought he was saved. Ernie went to the shopping mall and hunted around for a dark blue security guard uniform, and bought several. The next day he received a phone call from the factory personnel director about the security guard position. Ernie was dismayed that he had wasted money. He didn't have a job.

Set A, Theme 2: Carl wasn't working at the time. He was very concerned because he had very little left in his bank account. Several days later he had lunch with the president about becoming a broker. Carl thought he had impressed people when he gave his resume to the investment firm. Carl went to the department store and tried on some suits, and got a few. He felt that he was moving up again.

Set B, Theme 1: Ronnie thought she had it made because she thought she had done well in the audition for a keyboard player. Ronnie went to the music store, played some electric organs, and then purchased one. Later she got a message from the guitar player about playing keyboards. She wasn't in a band. Ronnie was depressed that she had run up her credit card.

Set B, Theme 2: Cindy was upset that she had blown her savings. She wasn't employed. Cindy was really happy about her tryout as dancer for a new musical. That night she met the director about the dancer position. Cindy got over to some stores, searched for, and bought some leotards. She believed her troubles were over.

more than one story, and the characters and objects used in set A that map to those in set B (e.g., stock broker and dancer) are not closely related semantically. Thus the paired stories based on a common theme represented analogous rather than literally similar instantiations of that theme. In terms of Halford's taxonomy, structural consistency was varied at the level of system similarity, while element similarity was held constant at a relatively low level and relational similarity was held constant at a moderate level.

Subjects rated 12 stories of the type illustrated in Table 2 for imagability. Later, subjects read 10 more stories and were asked to write down any of the rated stories of which they were reminded. The design and procedure were similar to those used in Experiments 1–2. In the previous experiments, subjects had always made semantic ratings of cue stories before they wrote down their remindings. To ensure that our general pattern of results is not dependent upon subjects making cue ratings, we changed our reminding instructions. After completing the distractor task, subjects were told to write down what they were reminded of by the cue stories, but were not asked to make ratings.

The event-related cue and target stories were either instantiations of the same theme (analogous) or of different themes (disanalogous). As before, items were cued in singleton and competition conditions (see Fig. 2).

Based on ARCS, we predicted that analogous stories would be retrieved more frequently than disanalogous stories, and that this difference would be greater in the competition than in the singleton condition.

Method

Materials. Materials consisted of 14 sets of four stories. Some stories were derived from the materials used by Rattermann and Gentner (1987) and Seifert et al. (1986). We constructed one set of basic events for each set of stories. Two unique story plots were created for each story set by rearranging the sequence of propositions. To avoid having the surface order of propositions covary with system similarity, event sequences shared between thematically similar stories were changed as much as possible without altering the underlying story plot. Each story within a set was written about a different set of actors, such as roommates, countries, or siblings. No content words or proper names were used in more than one story across the entire set of materials.

Two supersets of materials were created by assigning the stories in each of the story sets to one of two groups. The two stories from each story set assigned to each superset contained dissimilar themes. An equal number of cues and targets came from each superset. In order to have all stories appear in all conditions, we created 14 separate configurations of cue and target stories. Each configuration of materials was administered to three subjects. In all other respects, materials were composed in the same fashion as those in Experiment 1A.

In order to determine if, when reading our materials, people are sensitive to the factors we manipulated, 28 undergraduates attending UCLA completed questionnaires designed to assess the perceived similarity of the cue and target stories. Subjects were asked to rate ". . . how similar are the scenes being described" on a 6-point Likert scale (range: 1, completely dissimilar, to 6, completely similar). Analogous story pairs were rated more similar than disanalogous story pairs (4.81 vs 3.39), min $F'(1,53) = 36.70$, $p < .0001$; disanalogous story pairs were rated more similar than unrelated story pairs (2.00), min $F'(1,51) = 44.47$, $p < .0001$.

Subjects and procedure. Subjects were 42 UCLA undergraduates who participated either for pay or in order to meet a requirements for one of several psychology survey courses. During the encoding portion of the experiment, subjects rated the target stories on a 6-point Likert scale for imagability. Subjects were given 1 min to read each story. Unlike previous experiments, subjects did not make any semantic ratings of cue stories. In all other respects, the procedure was identical to that of Experiment 1A.

Results and Discussion

The proportion of target story types of which subjects were reminded by cue stories is shown in Fig. 6. Subjects recalled more analogous (consistent) stories than disanalogous (inconsistent) stories, min $F'(1,54) = 6.29$, $p < .05$. The interaction between competition and analogical similarity was not reliable, min $F'(1,54) = 2.39$, $p > .10$; however, both the subject and item ANOVAs were significant ($F(1,28) = 4.38$, $F(1,27) = 5.27$), $p < .05$ for each. As predicted, analogous stories were accessed more frequently than disanalogous stories within the competition condition, min $F'(1,54) = 6.76$, $p < .05$, but not within the singleton condition, min $F' < 1$. Probability of retrieving at least one related target in the competitor condition was 75% compared to 62% in the analogous single-

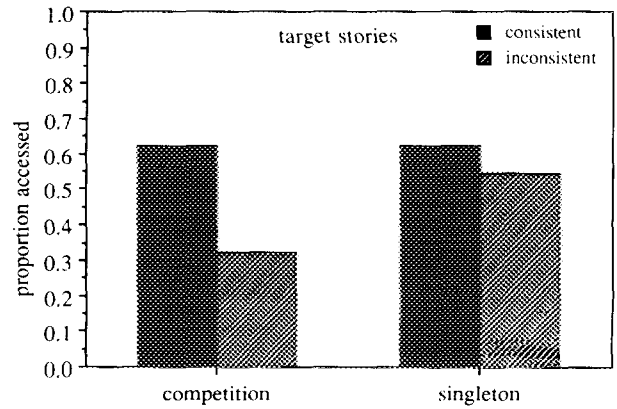

FIG. 6. Reminding for target stories in Experiment 3.

ton condition and 54% in the disanalogous singleton condition. Access of both analogous and disanalogous target sentences occurred in 19% of retrievals in the competitor condition. Remindings of unrelated targets across cue conditions showed a pattern very similar to that observed in previous experiments, increasing (except for a small reversal) across the competition, singleton consistent, singleton inconsistent, and unrelated-cue conditions (6, 13, 11, and 32%).

Comparison of our methodology and results with those of Rattermann and Gentner (1987) and Seifert et al. (1986) resolves some of the apparent empirical discrepancies involving the role of thematic similarity in episodic retrieval. As in our study, Rattermann and Gentner's minimal study paradigm probably resulted in subjects' encoding mostly superficial story representations. Accordingly, each story's theme did not differentiate it that much from other similar, but disanalogous stories. Rattermann and Gentner's failure to find a reliable effect of analogical similarity appears to have been a consequence of their use of the relatively weak singleton design. In Experiment 3, we obtained a strong effect of analogical similarity in the competition condition but not in the singleton condition. It would thus appear that even with minimally encoded stories, reminding is influenced by relational and system similarity. However, this influence is obscured in the absence of retrieval competition.

Seifert et al. (1986) also examined the role of system similarity in reminding. Immediately after reading two stories, subjects verified whether conclusion sentences from these stories had been seen before. The conclusion sentence from one story would be preceded by the conclusion sentence from the other story. The authors varied whether the stories were analogous or unrelated and how subjects encoded these stories.

Seifert et al. hypothesized that if subjects were reminded of an analogous story when reading a new story, then they should verify conclusion sentences relatively faster when they were preceded by conclusion sentences from analogous stories. Seifer et al. also used a singleton design. Congruent with Rattermann and Gentner's and our findings with a singleton design, when subjects only read the stories, subjects were not faster to verify conclusion sentences from an analogous story pair than they were to verify conclusion sentences from an unrelated pair. However, an effect of analogical similarity was shown when subjects explicitly compared both stories (Experiment 2). Seifert et al.'s results suggest that when subjects create abstract, thematically oriented representations of target stories (as opposed to the more superficial representations likely to have been formed by our subjects and those of Rattermann and Gentner), a robust effect of system similarity on reminding can be shown even with a singleton design.

ARCS SIMULATIONS

The results of the present experiments support the general predictions made by ARCS: When surface element similarity is equated between episodes in memory, structural consistency at the relational and system levels, as modulated by retrieval competition, has an impact on reminding. In order to more explicitly test its predictions, we used ARCS to simulate memory retrieval for hypothetical subjects using examples of materials based on those used in Experiments 2 and 3. These simulations required constructing representations of the cues and targets in the predicate-calculus notation that ARCS takes as input, and also constructing simple semantic networks linking the concepts involved in the texts. The simulations are only intended as rough sufficiency demonstrations regarding the model's predictions, as performance will of course vary with the representations and parameters assumed.

Simulation of Experiment 2

Table 3 shows an example of target and cue representations used in the simulation of Experiment 2. Each sentence in Experiment 3 was represented by a single two-place predicate representing the sentence's verb and two single-place predicates representing the sentence's subject and object. Cohyponyms were related in ARCS' semantic network by making them subordinates of a common superordinate. For example, "lawyer" and "judge" were both defined as subordinates of "legal expert."

The target structures for all conditions, along with four filler structures, were placed in ARCS' memory for each simulation. Each condition was simulated twice with different cues. There were a total of 12 target struc-

TABLE 3
Example of ARCS Representations for Materials Used in Experiment 2

Consistent target
The lawyer was warned by the paramedic.
(warn (obj-paramedic obj-lawyer) true t1-1)
(paramedic (obj-paramedic) t1-2)
(lawyer (obj-lawyer) t1-3
Inconsistent target
The attorney cautioned the fire fighter.
(caution (obj-attorney obj-fire-fighter) true t2-1)
(attorney (obj-attorney) t2-2)
(fire-fighter (obj-fire-fighter) t2-3)
Competitor cue
The judge was alerted by the rescue worker.
(alert (obj-rescue-worker obj-judge) c1-1)
(rescue-worker (obj-rescue-worker) c1-2)
(judge (obj-judge) c1-3)

Note. Two competitor target passages and one competitor cue passage shown in the predicate–calculus form used in an ARCS simulation. Sentences given to subjects are shown in italic type.

tures present in memory (i.e., four competitor targets, two singleton consistent targets, two singleton inconsistent targets, and four fillers) for each simulation. Unrelated cues had the same form as competitor and singleton cues, except that they had no direct relation to any of the target structures that were present in memory. Although unrelated cues were not directly related to any of the target structures, they could still make incidental contact with target structures if one or more of the subject, verb or object in the cue was a cohyponym of a subject, verb, or object in a target structure. ARCS' parameters were set at the following values: maximum activation, .99; minimum activation, −.30; inhibitory weights, −.20; excitatory weights, .005; decay, .005. Weights from the semantic unit were set by multiplying the similarity value for the mapped predicates (.1 for identical predicates, .08 for synonyms, −.4 for antonyms, and .06 for cohyponyms) by a factor of .01. The network was settled using the activation-updating rule proposed by Grossberg (1978). (See Thagard et al., 1990, for further simulation details.)[5]

[5] The present simulations were performed with two minor modifications of the ARCS algorithm described by Thagard et al. (1990). First, the predicate "cause" was permitted to act as a retrieval cue in the semantic network. Second, in generating mapping networks excitatory links were constructed between predicate-mapping units and argument-mapping units.

Across the 10 simulation runs, the mapping networks constructed ranged in size from 11–34 units interconnected by 18–246 links. All runs settled into a stable asymptotic state of the network after 26–178 cycles of updating. Table 4 shows the asymptotic activation values of the retrieved structures for each condition, averaged over the two simulations. The activation values do not correspond to response proportions in any precise quantitative fashion, as the model does not specify an explicit response rule for predicting when an item will be recalled by an individual subject. We simply note that the rank order of ARCS' activations roughly corresponds to the rank order of reminding frequencies. Although some deviations are apparent, the activations produced by ARCS reflect the major features of the human recall data: (a) the activation of the consistent competitor is higher than that of the inconsistent competitor, (b) the activation of the consistent singleton is higher than that of the consistent competitor, and (c) the activation of the inconsistent singleton is higher than that of the inconsistent competitor.

Simulation of Experiment 3

Table 5 shows an example of one of the story representations used in the simulation of Experiment 4. The predicate-calculus representations of the stories included both propositions representing what was explicitly stated in the text and propositions representing low-level inferences that we felt were necessary in order for subjects to comprehend the text. Inferences were included that either provided causal connections between events that were implicitly related in the text or that explicitly stated information that was implicit in the text. For example, given the passage, *But the next week, she got sick from eating some of the bread that had gone bad . . . she instructed the chef to add the preservative right away* (propositions cue1-ev9 to cue1-ev11 in Table 5), we included the plausible inference that the event described in the first sentence caused the event described in the second sentence (cue1-ev12 in Table 5). Similarly, given the passage, *Mrs. Keller didn't want to*

TABLE 4
ARCS Activation Values (Roman Type) and Obtained Retrieval Proportions (Italic Type) for Experiment 2

	Cue condition							
Recall type	Competition		Consistent		Inconsistent		Unrelated	
Consistent	.72	*.54*	.79	*.80*	—	—	—	—
Inconsistent	.10	*.42*	—	—	.80	*.67*	—	—
Unrelated	−.11	*.02*	−.12	*.06*	−.08	*.07*	.45	*.23*

TABLE 5
Examples of ARCS Representations for Materials Used in Experiment 3

Mrs. Keller was the health inspector at Wonderful Breads,
(health-inspector (obj-keller) true cue 1-attr1)
(works-for (obj-keller wonderful-breads) true cue1-st1)
(company (obj-wonderful-breads) true cue1-attr2)
who recently had started marketing a new, tasty bread that was popular but tended to spoil easily.
(market (obj-wonderful-breads obj-bread) true cue1-ev1)
(at-time (obj-recent-time) true cue1-st2)
(recent (obj-recent-time) true cue1-attr3)
(bread (obj-bread) true cue1-attr4)
(new (obj-bread) true cue1-attr5)
(tasty (obj-bread) true cue1-attr6)
(improved (obj-bread) true cue1-attr7)
(popular (obj-bread) true cue1-attr8)
(spoils (obj-bread) true cue1-attr9)
(sick (obj-unknown) true cue1-st3)
(person (obj-unknown) true cue1-attr10)
(can-cause (cue1-attr9 cue1-st3) true cue1-ev2)
Because of this, the head chief suggested that an expensive preservative be added to the recipe.
(suggest (obj-chef cue1-plan1) true cue1-ev3)
(chef (obj-chef) true cue1-attr11)
(works-for (obj-chef wonderful-breads) true cue1-st4)
(cause (cue1-ev2 cue1-ev3) true cue1-ev4)
(disable (cue1-attr9 cue1-ev5) true cue1-plan1)
(add (obj-preservative obj-bread) true cue1-ev5)
(preservative (obj-preservative) true cue1-attr12)
(expensive (obj-preservative) true cue1-attr13)
Mrs. Keller didn't want to change the recipe.
(not-want (obj-keller cue1-ev4) true cue1-st5)
It would be a lot of trouble and would change the flavor
(troublesome (cue1-ev5) true cue1-attr14)
(change (cue1-ev4 cue1-attr6) true cue1-ev6)
(undesirable (cue1-ev6) true cue1-attr15)
(cause (cue1-ev5 cue1-st5) true cue1-ev7)
(cause (cue1-attr14 cue1-st5) true cue1-ev8)
But the next week, she got sick from eating some of the bread that had gone bad.
(eat (obj-keller obj-bread) true cue1-ev9)
(sick (obj-keller) true cue1-st6)
(cause (cue1-ev9 cue1-st6) true cue1-ev10)
(before (cue1-st5 cue1-st6) true cue1-st7)
She got so sick she instructed the chef to add the preservative right away.
(instruct (obj-keller obj-chef cue1-plan1) true cue1-ev11)
(cause (cue1-st6 cue1-ev11) true cue1-ev12)
(changed-mind (obj-keller cue1-st5) true cue1-ev13)

Note. One story given to subjects is shown in the predicate-calculus form used in an ARCS simulation. Sentences from story given to subjects are shown in italic type.

change the recipe. It would be a lot of trouble and would change the flavor. (cue1-st5 to cue1-ev6), we included the inference that changing the flavor was an undesirable thing to do (cue1-attr15). Items in the stories were represented semantically in ARCS by means of a semantic network constructed in the same manner as that used in the simulation of Experiment 2.

We constructed representations of two cue stories, two target stories that were related to the cue stories, and six filler stories that were unrelated to the cue and the target stories. Each target story was consistent with one of the cues and inconsistent with (but semantically related to) the other cue. This allowed us to perform two simulations each of the competitor, consistent singleton and inconsistent singleton conditions by using each of the two cue stories as probes in separate simulations. The simulation of each condition had the appropriate target stories in memory, along with all six filler stories. The unrelated condition was simulated by placing all stories into memory and using an arbitrarily chosen filler story as the cue. Parameters used for the simulations were: maximum activation, .99; minimum activation, −.30; inhibitory weights, −.03; excitatory weights, .01; decay, .01. The parameter values were reduced relative to those used in the simulation of Experiment 2 because the mapping networks were much larger for the complex representations required for the story materials of Experiment 3, and high parameter values tend to make activations unstable for large networks. Weights from the semantic unit were set by multiplying the similarity values of mapped predicates by a factor of .10. As in the simulation of Experiment 2, the Grossberg updating rule was used to settle the network. Simulations for Experiment 3 were run on a 16000-processor CM2 Connection Machine using a version of ARCS written in *LISP, which updates activations in parallel.

Across the 10 simulation runs, the mapping networks constructed ranged in size from 92 to 503 units interconnected by 1276–13978 links. All runs settled into a stable asymptotic state of the network after 223–396 cycles of updating. The mean activations for the retrieved structures are shown in Table 6. As in the simulation of Experiment 2, the activations

TABLE 6
ARCS Activation Values (Roman Type) and Obtained Retrieval Proportions (Italic Type) for Experiment 3

	Cue condition							
Recall type	Competition		Consistent		Inconsistent		Unrelated	
Consistent	.65	*.62*	.77	*.62*	—	—		
Inconsistent	.39	*.32*	—	—	.72	*.54*		
Unrelated	.21	*.07*	.30	*.13*	.30	*.12*	.20	*.32*

produced by ARCS capture the most important aspects of the reminding frequencies: (a) the activation of the consistent competitor is higher than that of the inconsistent competitor, (b) the activation of the consistent singleton is higher than that of the consistent competitor, and (c) the activation of the inconsistent singleton is higher than that of the inconsistent competitor.

GENERAL DISCUSSION

Structural Consistency and Retrieval Competition

The present findings clarify a number of basic constraints on the mechanisms of analogical reminding. Our results provide strong evidence that memory access in reminding is systematically competitive and that structurally consistent mappings linking cues and stored memory representations provide a basic constraint on reminding. These effects of structural consistency and of retrieval competition were observed even though cues and stored structures had never been studied together, had no lexical overlap, and, in the case of surrounds activated by target cues in Experiments 1A and 1B, had no direct semantic links. These effects were observed when structural consistency was varied by manipulating relational similarity, holding element similarity constant at a high level (Experiments 1–2) and when it was varied by manipulating system similarity (correspondences between themes), holding element similarity constant at a low level and relational similarity constant at a moderate level (Experiment 3).

Several aspects of the experimental design allow us to draw firmer conclusions than made possible by previous investigations of structural effects on reminding. We obtained an effect of structural consistency under conditions in which similarity of individual concepts was equated (because across counterbalancing conditions the same words were used in both the consistent and inconsistent targets), and surface correspondences at the level of word order were controlled (by random assignment of active versus passive voice to cues and target sentences). Our measure of access to texts did not depend on successful completion of postaccess processes such as problem solving or explanation, thus localizing the impact of consistency at the initial access stage itself. Moreover, because the cue sentences did not provide specific processing goals, the consistency effect proved separable from pragmatic influences on reminding. Finally, the observed effect of structural consistency manifested itself in an access measure based simply on recovery of content words from the text, rather than recovery of relational structure (although subjects in Experiments 1–2 accurately recalled case–role assignments in addition to individual content words). Our results therefore support the conclusion that structural relations between a cue and stored memory representations constitute one of the basic constraints on initial access.

The results of Experiment 3 indicate that episodes in memory can be retrieved not only on the basis of element and relational similarity, but also system similarity. Cues and targets were related by abstract thematic relations (*cf.* the "Thematic Organizing Packets" proposed by Shank, 1982). It should be emphasized, however, that our results do not suggest that reminding can occur in the absence of *any* semantic overlap. In Experiments 1–2, both relations and objects were similar across related cues and targets; in Experiment 3, relations but not objects were similar. Further research is needed to explore whether system similarity can influence reminding even when neither first-order relations nor objects are similar across cue and target.

The present study also provides clear evidence that memory access is fundamentally competitive. Previous studies of analogical reminding have failed to consider the possible role of retrieval competition on recall of individual items. The present results provide an explanation of why previous studies have often failed to obtain clear evidence that structural consistency influences reminding. Across Experiments 1–2, the mean advantage of structurally consistent target sentence passages relative to inconsistent target sentence passages was 23% in competition conditions, but only 13% in singleton conditions. The difference in mean reminding advantage for structurally consistent passages across the two conditions was even more pronounced for surrounds: 33% in competition conditions versus only 4% in singleton conditions. A similar pattern was observed in Experiment 3. Access of inconsistent targets suffered even though instructions to subjects stressed that they should report any and all texts of which they were reminded (a procedure that might be dubbed "modified modified free reminding," by analogy to an important, if cumbersomely named, free-recall paradigm introduced by Barnes & Underwood, 1959). These results support our claim that previous studies of analogical reminding have underestimated the importance of structural consistency because they only examined the equivalent of our singleton conditions. Of course, competition cannot increase the absolute probability of reminding for any target; however, it can amplify the *difference* between the probabilities of retrieving targets that differ in how well they map to the cue.

Implications for Other Retrieval Models

The present results also help to evaluate other models of analogical reminding and general memory retrieval. The MAC/FAC ("many are called but few are chosen") model of Gentner and Forbus (1991) is similar to ARCS in predicting that memory retrieval will be influenced by both surface similarity and structural consistency. In MAC/FAC, retrieval is based on both surface descriptions and higher-order relations. Computationally, retrieval is a two-step process. In the first stage (MAC), the episode in long-term memory that has the most commonalities (in terms of predicates, functions, and connectives) with the probe story is retrieved. Stories are represented in MAC as vectors of content words, and similarity is computed by taking the dot product between the probe and each story in memory. Any other stories that yield dot products within 10% of the best match are also retrieved. The second stage (FAC) computes how well each retrieved first-stage story matches the cue, based on common relational structure and object descriptions. Stories are represented in FAC in predicate-calculus form. The episode with the highest match, along with any story within 10% of the best match, is retrieved. Thus MAC is sensitive only to element similarity, whereas FAC is influenced by both element similarity and structural consistency.

Because the current implementation of MAC/FAC does not include a semantic network or other scheme for finding links between nonidentical concepts, it would not be able to model remindings between stories that lack lexical overlap, such as those used in the present experiments. Presumably, however, the model could in principle be extended to include a semantic hierarchy of the sort employed by ARCS, so that retrieval could be based on shared meaning components. Unlike ARCS, retrieval competition does not play a direct role in the operation of MAC/FAC; however, the assumption that the best match will always be retrieved (regardless of its absolute goodness) leads to the prediction that retrieval of weaker matches will suffer relatively more in a competition design. Thus although MAC/FAC implements retrieval in a somewhat different way than does ARCS, the two models are broadly similar in their qualitative predictions about constraints that govern reminding.

The present findings are inconsistent with memory-retrieval models based solely on unstructured feature vectors (e.g., Eich, 1982; Murdock, 1982; see discussion by Ratcliff & McKoon, 1989), as well as with models of analogical access based solely on "mere appearance" mapping (Falkenhainer, Forbus, & Gentner, 1989). In contrast, our evidence for an impact of structural consistency on reminding is consistent with ARCS and MAC/FAC, as well as with other memory models in which encoding and retrieval of propositional knowledge is sensitive to case–role assignments (e.g., Anderson, 1983; Hinton, 1981; Shastri & Ajjanagadde, 1993). Similarly, most artificial-intelligence models of case-based reasoning favor access to episodes that are structurally consistent with the cue in terms of similar goal/plan relationships (e.g., Hammond, 1989; Kolodner, Simpson, & Sycara, 1985; Schank & Leake, 1989).

Future Directions

The present experiments used relatively brief delays between study and test; further work is required to investigate the robustness of structural influences on reminding over longer delays. In addition, future research needs to examine the relationship between structural constraints on access and more pragmatic pressures, such as goal relevance. Although the focus of the present investigation was on the role of structural relations in reminding, it is apparent that analog retrieval, like analogical mapping, is the product of interactions among multiple constraints, of which structural overlap is but one.

Future work also needs to address the role of comprehension and inference processes in reminding. Comparison of the results of Experiment 2 with those obtained in Experiments 1A and 1B suggests that a meaningful context, which supports greater number of inferences, is likely to increase the magnitude of the advantage for structurally consistent targets. The ARCS model, which has no inferential capability, is unable to account for the role of context in modulating the effect of structural consistency. It is possible that the structural effects observed in the present study are side effects of altering argument structure. Reversing object and subject not only can alter situational inferences but can also change individual word sense (*cf.* the different senses of *shark* in *The shark ate the surfer* to *The surfer ate the shark*). While the present study has demonstrated that the structural differences between sentences and passages do indeed affect reminding, it has not determined whether that effect is due to the structural differences *per se*, the different inferences the structural differences engender or a combination of the two.

Inferences are likely to be especially critical in providing the capacity to model remindings based on higher-order relations (i.e., system similarity). ARCS, which uses only taxonomic relations between individual concepts as retrieval paths, lacks any capacity to infer, for example, that the overall theme of a passage is "retaliation." Consequently, ARCS is unable to use such implicit abstractions to guide retrieval by relating episodes with similar abstract themes. We have recently attempted to overcome the above limitations by developing a hybrid symbolic–connectionist model that integrates text inferencing and analogical reminding (Lange, Melz, Wharton, & Holyoak, 1990; Lange & Wharton, 1992, in press).

The present study also has suggestive implications for case-based reasoning and explanation-based learning models that retrieve multiple episodes before beginning problem-solving activity (e.g., Kolodner, 1988). Recalling multiple episodes allows these models to induce relevant generalizations. However, it is unclear under what conditions people are likely to recall multiple episodes in response to a single retrieval cue. In the present experiments, recall of multiple episodes to a single cue was infrequent. Across Experiments 1–2, the percentage of recall instances for the competition condition in which two related episodes were recalled was only 12% for target sentences and 9% for surround sentences. For Experiment 3, the percentage of recall instances for the competition condition in which two related episodes were recalled was 19%. We do not believe that these infrequent multiple remindings were an artifact of our instructions. Subjects were told that they might be reminded of multiple stories, in which case they were to recall them all if possible. It is possible that the relatively low rate of multiple-episode remindings was due to the fact that the passages in each pair of our competitor conditions were not good analogs of each other. As a result, the first passage recalled would be a poor cue for its mate. We would expect a higher rate of multiple remindings (as well as more blending) if both targets in memory were analogous to each other. Future research should explore the contribution made by factors such as episode characteristics, reminding motivation, and domain expertise to multiple-episode reminding.

A related issue concerns the role of intentional retrieval mechanisms in reminding. Although anecdotal reports of everyday reminding suggest that it often occurs in the absence of any intention to be reminded (Schank, 1982), most laboratory studies of reminding, including the present one, have provided subjects with instructions to deliberately use cues to generate memory retrieval (Gentner & Landers, 1985; Johnson & Seifert, 1990; Rattermann & Gentner, 1987). It may prove useful to contrast two types of memory mechanisms, which might be termed analogical *reminding* and analogical *priming*. The general difference between reminding and priming is that reminding involves a "backwards" influence of current processing of the cue on the activation of traces of previous episodes, whereas priming involves a "forward" influence of traces of previous episodes to influence current processing of the cue. Awareness of retrieval is not necessary for priming effects to occur. We suspect that a complete model of analogical access will require attention to both reminding and priming, as well as interactions between them.

It seems likely that the effect of system and relational similarity on memory retrieval is greater in intentional reminding than in unintentional priming. Explicit instructions typically provide temporal and contextual cues (e.g., "Does this remind you of any of the stories you read a few minutes ago in this room?"), as well as specific content cues. Context cues allow the person to focus attention on an episodically defined subset of the entire contents of memory. This allows finer discriminations of relations involving the content cues than would otherwise be possible. Models of analogical reminding might usefully incorporate mechanisms for integrating context and content cues that have been employed in general models of retrieval (e.g., Gillund & Shiffrin, 1984; Raajimakers & Shiffrin, 1981; Ratcliff & McKoon, 1988). More generally, it is important to remember that analogical reminding surely depends on the same overall memory system used for more literal recall and recognition.

REFERENCES

Anderson, J. R. (1974). Retrieval of propositional information from long-term memory. *Cognitive Psychology*, **5**, 451–474.

Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard Univ. Press.

Anderson, J. R., & Milson, R. (1989). Human memory: An adaptive perspective. *Psychological Review*, **96**, 703–719.

Barnden, J. A. (in press). On the connectionist implementation of analogy and working memory matching. In J. A. Barnden & K. J. Holyoak (Eds.), *Advances in connectionist and neural computation theory: Reminding and Metaphor* (Vol. 3). Norwood, NJ: Ablex.

Barnes, J. M., & Underwood, B. J. (1959). "Fate" of first-list associations in transfer theory. *Journal of Experimental Psychology*, **58**, 97–105.

Bjork, R. A. (1989). Retrieval inhibition as an adaptive mechanism in human memory. In G. H. Bower (Ed.), *Advances in the psychology of learning and motivation* (Vol. 23, pp. 309–330). New York: Academic Press.

Clark, H. H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior*, **12**, 335–359.

Eich, J. M. (1982). A composite holographic associative recall model. *Psychological Review*, **89**, 627–661.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, **41**, 1–63.

Fodor, J. A., & McLaughlin, B. P. (1990). Connectionism and the problem of systematicity: Why Smolensky's solution doesn't work. *Cognition*, **35**, 183–204.

Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. In S. Pinker & J. Mehler (Eds.), *Connections and symbols* (pp. 3–71). Cambridge, MA: MIT Press.

Gentner, D. (1983). Structure-mapping: A theoretical framework. *Cognitive Science*, **7**, 155–170.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony, *Similarity and analogical reasoning* (pp. 199–241). Cambridge, UK: Cambridge.

Gentner, D., & Forbus, K. D. (1991). MAC/FAC: A model of similarity-based retrieval. *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society* (pp. 504–509). Hillsdale, NJ: Erlbaum.

Gentner, D., & Landers, R. (1985). Analogical reminding: A good match is hard to find. *Proceedings of the International Conference on Systems, Man and Cybernetics* (pp. 607–613). Tucson, AZ.

Gentner, D., Rattermann, M. J., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology*, **25**, 431–467.

Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, **10**, 277–300.

Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, **12**, 306–355.

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, **15**, 1–38.

Gick, M. L., & McGarry, S. J. (1992). Learning from mistakes: Inducing analogous solution failures to a source problem produces later successes in analogical transfer. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **18**, 623–639.

Goldstone, R. L., Medin, D. L., & Gentner, D. (1991). Relational similarity and the non-independence of features in similarity judgments. *Cognitive Psychology*, **23**, 222–262.

Grossberg, S. (1978). A theory of visual coding, memory, and development. In E. L. J.

Leeuwenberg & H. F. J. Buffart (Eds.), *Formal theories of visual perception.* New York: Wiley.

Halford, G. S. (1992). Analogical reasoning and conceptual complexity in cognitive development. *Human Development,* **35,** 193–217.

Halford, G. S., & Wilson, W. H. (1980). A category theory approach to cognitive development. *Cognitive Psychology,* **12,** 356–411.

Halford, G. S., Wilson, W. H., Guo, J., Wiles, J., & Stewart, J. E. M. (in press). Connectionist implications for processing capacity limitations in analogies. In K. J. Holyoak & J. A. Barnden (Eds.), *Advances in connectionist and neural computation theory: Analogical connections* (Vol. 2). Norwood, NJ: Ablex.

Hammond, K. (1989). *Case-based planning.* Boston: Academic Press.

Hinton, G. E. (1981). Implementing semantic networks in parallel hardware. In G. E. Hinton & J. A. Anderson (Eds.), *Parallel models of associative memory.* Hillsdale, NJ: Erlbaum.

Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory & Cognition,* **15,** 332–340.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science,* **13,** 295–355.

Johnson, H. M., & Seifert, C. M. (1992). The role of predictive features in retrieving analogical cases. *Journal of Verbal Learning and Verbal Behavior,* **31,** 648–667.

Keane, M. (1988). *Analogical problem solving.* Chichester, England: Ellis Horwood.

Kintsch, W., & van Dijk, T. A. (1978). Toward a model of text comprehension and reproduction. *Psychological Review,* **85,** 363–394.

Kolodner, J. (1988). Retrieving events from a case memory: A parallel implementation. *Proceedings of the DARPA Workshop on Case-Based Reasoning.* San Mateo, CA: Morgan Kaufmann.

Kolodner, J., Simpson, R., & Sycara, K. (1985). A process model of case-based reasoning in problem solving. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence.* Los Angeles, CA.

Lange, T., Melz, E., Wharton, C., & Holyoak, K. (1990). Analogical retrieval within a hybrid spreading-activation network. In D. S. Touretzky, J. L. Elman, T. J. Sejnowski, & G. E. Hinton (Eds.), *Proceedings of the 1990 Connectionist Models Summer School.* San Mateo, CA: Morgan Kaufmann.

Lange, T. E., & Wharton, C. M. (1992). REMIND: Integrating language understanding and episodic memory retrieval in a connectionist network. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society* (pp. 576–581). Hillsdale, NJ: Erlbaum.

Lange, T. E., & Wharton, C. M. (in press). REMIND: Retrieval from Episodic Memory by INferencing and Disambiguation. In J. A. Barnden & K. J. Holyoak (Eds.), *Advances in connectionist and neural computation theory: Reminding and Metaphor* (Vol. 3). Norwood, NJ: Ablex.

McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General,* **114,** 159–188.

McDougal, T., Hammond, K., & Seifert, C. (1991). A functional perspective on reminding. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society* (pp. 510–521). Hillsdale, NJ: Erlbaum.

McKoon, G., Ratcliff, R., & Seifert, C. (1989). Making the connection: Generalized knowledge structures in story understanding. *Journal of Memory and Language,* **28,** 711–734.

Miller, G. A., Fellbaum, C., Kegl, J., & Miller, K. (1988). WORDNET: An electronic lexical reference system based on theories of lexical memory. *Revue Quebecoise Linguistique*, **17**, 181–213.

Murdock, B. B., Jr. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, **89**, 609–626.

Premack, D. (1983). The codes of man and beasts. *Behavioral and Brain Sciences*, **6**, 125–167.

Raajimakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, **88**, 93–134.

Ratcliff, R., & McKoon, G. (1988). A retrieval theory of priming in memory. *Psychological Review*, **95**, 385–408.

Ratcliff, R., & McKoon, G. (1989). Similarity information versus relational information: Differences in the time course of retrieval. *Cognitive Psychology*, **21**, 139–155.

Rattermann, M. J., & Gentner, D. (1987). Analogy and similarity: Determinants of accessibility and inferential soundness. In *Proceedings of the Ninth Annual Meeting of the Cognitive Science Society* (pp. 22–34). Hillsdale, NJ: Erlbaum.

Read, S. J., & Cesa, I. L. (1991). This reminds me of the time when . . . : Expectation failures in reminding and explanation. *Journal of Experimental Social Psychology*, **27**, 1–25.

Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **13**, 629–639.

Ross, B. H. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **15**, 456–468.

Schank, R. C. (1982). *Dynamic memory*. New York: Cambridge Univ. Press.

Schank, R. C., & Leake, D. B. (1989). Creativity and learning in a case-based explainer. *Artificial Intelligence*, **40**, 353–385.

Seifert, C. M., McKoon, G., Abelson, R. P., & Ratcliff, R. (1986). Memory connections between thematically similar episodes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **12**, 220–231.

Shastri & Ajjanagadde (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings. *Behavioral and Brain Sciences*, **16**, 417–494.

Spencer, R. M., & Weisberg, R. W. (1986). Context-dependent effects on analogical transfer. *Memory & Cognition*, **14**, 442–449.

Thagard, P., Holyoak, K. J., Nelson, G., & Gochfeld, D. (1990). Analog retrieval by constraint satisfaction. *Artificial Intelligence*, **46**, 259–310.

Wharton, C. M., Holyoak, K. J., Downing, P. E., Lange, T. E., & Wickens, T. D. (1991). Retrieval competition in memory for analogies. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society* (pp. 528–533). Hillsdale, NJ: Erlbaum.

Accepted February 21, 1993